
Co-Evolution of Form and Function in the Design of Autonomous Agents: Micro Air Vehicle Project

Magdalena D. Bugajska

Alan C. Schultz

Navy Center for Applied Research in Artificial Intelligence (Code 5515)
Naval Research Laboratory, Washington, DC 20375-5000, USA
magda, schultz {@aic.nrl.navy.mil}
tel. (202) 404-4946 fax (202) 767-3172

Abstract

This paper addresses issues of co-evolution of form and function for autonomous vehicles, specifically evolving morphology and control for an autonomous micro air vehicle (MAV). The evolution of an optimal minimum sensor suite and reactive strategies for navigation and collision avoidance for the simulated MAV is described. The details of the implementation of the simulated aircraft, the environment, and the two cooperating genetic algorithm-based systems, SAMUEL and Genesis, used for evolution, are presented, as are preliminary results.

1 INTRODUCTION

The co-evolution of form and function is the way all living organisms evolved in nature. If nature's example is to be followed, the form and function of autonomous agents should be co-evolved in a similar manner.

In this study, the concept of the co-evolution of form and function is applied to the Micro Air Vehicles (MAVs) domain. MAV should be thought of as an aerial autonomous agent, a six-degree-of-freedom vehicle whose mobility allows us to deploy a useful micro payload to a remote or otherwise hazardous location where it may perform variety of missions, including reconnaissance and surveillance, targeting, tagging, and bio-chemical sensing. The design of MAVs calls for aircraft that is at least an order of magnitude smaller than any current flying system; the target vehicle whose model is used for this study, has a wingspan of 6 inches (15 cm). Due to the size of the aircraft as well as the variety of applications, the design of the sensory payload and the controller of the MAV are quite complex, as are the relationships between them. The design issue addressed explicitly in this study is minimization of weight and power requirements. The goal of the study is to evolve a sensor suite with a minimal number of sensors, which allows for the most efficient task-specific control. The

experimental task requires MAV to navigate to a specified target location, while avoiding collision with obstacles. The co-evolution is performed in simulation using two cooperating genetic algorithm-based systems, SAMUEL [Grefenstette 91] and GENESIS [Grefenstette 84].

The remainder of this paper outlines the work done up to this date, and then goes into details about our implementation of co-evolution of the behaviors required for collision-free navigation and the characteristics of a sensor suite that would allow the MAV to perform its task with a maximum efficiency. The simulated environment, aircraft, and sensors are described and the details of the two learning systems are provided as well. Finally, some initial results are presented, and the future direction of the research is outlined.

2 RELATED WORK

Evolutionary algorithms have been shown to be effective procedures for searching large and complex spaces. They have been successfully applied to automate the design of robots' morphology as well as the design of the controllers, but the concept of co-evolution of form and function has surfaced only recently.

There has been a great deal of work done in the area of evolution of function for autonomous robots. [Nolfi 94] evolves neural controllers for collision-free navigation for mobile robots. [Harvey 92] reports on evolving neural control systems for the task of exploration. [Schultz 91] used a genetic algorithm-based system, SAMUEL, to learn reactive rule-based strategies for collision-free navigation for an autonomous underwater vehicle (AUV) as well as shepherding [Schultz 96] and tracking for other mobile robots. [Sammur 92] demonstrates machine learning of a reactive strategy to control a dynamic system by observing a controller that is already skilled in the task. While [Floreano 96] discusses similar work, the evolutionary process in this study is carried entirely online on the physical robot.

In parallel to research of techniques of evolution of function, similar research is being done in the area of evolution of form. [Funes 97] applied evolutionary

Report Documentation Page

*Form Approved
OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 2000	2. REPORT TYPE	3. DATES COVERED 00-00-2000 to 00-00-2000	
4. TITLE AND SUBTITLE Co-Evolution of Form and Function in the Design of Autonomous Agents: Micro Air Vehicle Project		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Navy Center for Applied Research in Artificial Intelligence (NCARAI), 4555 Overlook Avenue SW, Washington, DC, 20375		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited			
13. SUPPLEMENTARY NOTES GECCO-2000 Workshop on Evolution of Sensors in Nature, Hardware, and Simulation, Las Vegas, NV; July 8, 2000			
14. ABSTRACT This paper addresses issues of co-evolution of form and function for autonomous vehicles specifically evolving morphology and control for an autonomous micro air vehicle (MAV). The evolution of an optimal minimum sensor suite and reactive strategies for navigation and collision avoidance for the simulated MAV is described. The details of the implementation of the simulated aircraft, the environment, and the two cooperating genetic algorithm-based systems, SAMUEL and Genesis, used for evolution, are presented, as are preliminary results.			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	
			18. NUMBER OF PAGES 5
			19a. NAME OF RESPONSIBLE PERSON

techniques to the design of structures assembled out of parts. [Husbands 96] uses a distributed genetic algorithm and a distributed genetic algorithm hybridized with gradient decent techniques to evolve the cross-section of optimal aircraft wing-boxes. [Lichtensteiger 99] presents a study of evolution of the morphology of the compound eye. In [Lund 97] evolution of a morphology of an auditory hardware is discussed. [Mark 98] presents a framework for the study of sensor evolution in a continuous 2-dimensional virtual world (XRaptor).

Finally, in recent years work has began on co-evolving form and function for autonomous agents. [Sims 94] presents a system for the co-evolution of morphology and behavior of virtual creatures that compete in physically simulated three-dimensional world. In [Lee 96] a hybrid genetic programming/genetic algorithm approach is presented that allows for evolution of both controllers and robot bodies to achieve behavior-specified tasks. [Lund 98] introduces a LEGO simulator that allows the user to co-evolve controllers and body plans using interactive genetic algorithm in simulation before constructing the LEGO robots. [Balakrishnan 96] presents the comparative study of evolution of a control system given fixed sensor architecture, and co-evolution of sensor characteristics (placement and range) and the control architecture for the task of box pushing.

The work presented here is related to these projects, but differs in several aspects. The result of this study is a learning system consisting of two cooperating genetic algorithm-based systems that allows for co-evolving control behaviors and the sensor suite for the MAV whose task is to navigate to a specified target location while avoiding obstacles. While the majority of the previous work involved evolution of neural controllers, our approach implements evolution of stimuli-response rules. The sensors characteristics initially evolved include the number of sensors and their beam width, with the future possibility of evolution of range and explicit placement of each sensor. Also, even though the evolution is performed in simulation, the simulator closely models the real aircraft and its environment. Finally, the control behaviors are not evolved in a specific setup of an environment as in [Balakrishnan 96], [Lund 98], and [Lee 96], but rather each single trial is performed in a randomly and dynamically created environment.

3 EVOLUTION OF SENSOR DESIGN AND CONTROL FOR MAV

The objective of the study is to evolve a sensor suite with a minimal number of sensors, which allows for the most efficient task-specific control. This section gives an overview of the learning system used to co-evolve the sensor characteristics and the control of the MAV whose task is a collision-free navigation to a specified target location.

The learning system used for co-evolution of form and function in this study is composed of two cooperating genetic algorithm-based systems, SAMUEL and

GENESIS. SAMUEL evolves the stimuli-response rules to control the MAV, while GENESIS is used to evolve characteristics of the sensors for the aircraft, for example: sensor range, area coverage, and placement.

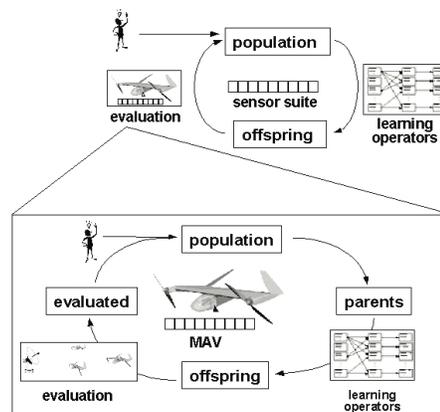


Figure 1: Cooperating genetic algorithm-based systems.

The two systems create a loop (see Figure 1) in which the output from one learning system is the input to the other one. Each member of population being evaluated by GENESIS represents a specific sensor configuration, which has to be evaluated by SAMUEL. Since it is assumed that the weight and power requirements can be fulfilled just by decreasing number of sensors on board of the MAV, the GENESIS evaluates each member of the population based on the number of the sensors in the suite and its task performance in the simulated environment as defined by the performance value returned by SAMUEL. This process is repeated until the minimum number of sensors is found that ensures the maximum efficiency of control given that sensor suite for the specified task or until maximum number of generations is reached.

The following sections will discuss the details of the evolution of both form (Section 5.0) and function (Section 4.0). The goals of each learning task will be reviewed, followed by implementation details and the short description of both learning systems, the representations and the fitness functions used.

4 EVOLUTION OF FUNCTION

In this section, the details of the MAV's control task and its process of evolution are discussed. Experimental details of the simulated environment, aircraft, and sensors are provided along with the details of the learning system used.

4.1 PROBLEM DESCRIPTION

The MAV must be able to efficiently and safely navigate in 3-Dspace among obstacles (trees) to a target location. The desired behavior should maximize the number of times the MAV reaches the target location while minimizing the distance traveled to that location. This problem includes several features that make it a challenging machine learning problem, e.g.: a weak

domain knowledge (e.g. no predictive model of obstacles or the goal), incomplete state information provided by discrete (possibly noisy) sensors, a large state space, and, of course, delayed payoff. The generality of evolved control is ensured due to a random setup of the environment and the MAV's position in it for every evaluation.

4.2 PROBLEM REPRESENTATION

4.2.1 Environment

Since the learning is being done in simulation, the MAV and its environment have to be modeled. To model the world as well as the aircraft itself, a high-fidelity, 3-D flight simulator is used, which includes an accurate parametrized model of a 6-inch MAV. The low level control for the MAV was implemented using a number of PID controllers in such a way that the plane could be controlled through changes made to turn rate and altitude of the aircraft. The trees (obstacles) were modeled as spheres (treetops) on top of cylinders (trunks). Any contact between the plane and the tree constituted a collision. The density of trees was user-defined as a number of trees per square foot assuming uniform distribution. At the beginning of each simulated flight, the MAV was placed in a random location within a specified area away from the target. The target remained stationary throughout the flight.

4.2.2 Sensors

There is a wide variety of sensors that could be implemented on the MAV, but the final make up of the sensor suite is constrained by the size, weight, and power capacity of the vehicle. It is assumed that the MAV has a sensor, which returns the relative range and bearing to the target. Also, the aircraft is equipped with a number of range sensors similar in capability to radar or sonar. Each sensor is capable of detecting obstacles and returning the range of the closest object within the sector covered by that sensor. The exact makeup of these sensors is learned by the evolution of form as described in Section 5.0.

4.2.3 Actions/Effectors

There is a discrete set of actions available to control the MAV. In this study, the only action that is considered specifies discrete turning rates for the MAV. The control variable `turn_rate` is between -20 and 20 degrees in 5 degrees increments. The altitude of the plane is held constant by the underlying PID controllers.

4.3 IMPLEMENTATION OF EVOLUTION

4.3.1 The Learning System

The behaviors required for navigating MAV to the target location while avoiding obstacles, which are represented as a collection of stimulus-response rules, are learned in the SAMUEL rule learning system. SAMUEL is a machine learning program that uses standard genetic

algorithms and other competition-based heuristics to solve sequential decision problems. It features Lamarckian operators (specialization, generalization, merging, avoidance, and deletion) that modify decision rules on the basis of observed interaction with the task environment. The original system implementation is described in greater detail in [Grefenstette 91].

4.3.2 Representation

SAMUEL implements behaviors as a collection of stimulus-response rules. Each stimulus-response rule consists of conditions that match against the current sensors of the autonomous vehicle, and an action that suggests action to be performed by it. An example of a rule (gene) might be:

```
RULE 122
  IF bearing = [-20, 20] AND sonar4 < 45
  THEN SET turn_rate = -100
```

Each rule has an associated strength with it as well as number of other statistics. During each decision cycle, all the rules that match the current state are identified. Conflicts are resolved in favor of rules with higher strength. Rule strengths are updated based on rewards received after each training episode.

4.3.3 Fitness Function

The simulation is divided into episodes that begin with placement of the MAV at a random distance (between 500 and 750 units) away from the target facing in random direction, which is followed, by a random placement of trees in the environment. The episodes end with either a successful arrival of the MAV at the target location, a loss of the MAV due to energy running out, or a loss of the MAV due to collision with an obstacle (tree or ground). The arrival is successful if the MAV approaches the target location within 15 units. The payoff for this study is defined as:

$$\text{payoff} = \begin{cases} 0.0 - 0.4, & \text{if MAV reached the target location,} \\ & f(\text{distance traveled}) \\ 0.7 - 1.0, & \text{if MAV collided with an obstacle or} \\ & \text{the time limit was reached, } f(\text{distance to goal}) \end{cases} \quad (\text{Eqn. 1})$$

5 EVOLUTION OF FORM

In this section, the details of the MAV's sensor suite configuration and its process of evolution are discussed.

5.1 PROBLEM DESCRIPTION

Due to the size of the MAV and the variety of applications, the design of the sensory payload for MAVs involves many tradeoffs. The design issue addressed in this study is the minimization of weight and power requirements. The objective of this study is to evolve a sensor suite with a minimal number of sensors that guarantees an efficient task-specific control. The sensor design is being evolved along with the decisions rules that control the actions of the MAV.

5.2 PROBLEM REPRESENTATION

Given a task of evolving the characteristics of a sensor suite, the sensor model in Figure 2 was assumed.

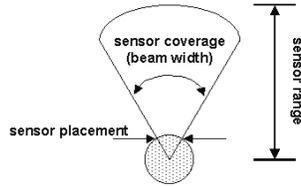


Figure 2: Sensor Model.

The sensor is similar in capability to a radar or sonar i.e. it returns the range to the closest obstacle in its field of view. The evolvable sensor characteristics include:

1. number of sensors
2. minimum range of the individual sensor
3. maximum range of the individual sensor
4. beam width of the individual sensor
5. placement of individual sensor

Given these characteristics, there are two types of the suites that can be designed: homogeneous and heterogeneous (see Figure 3).

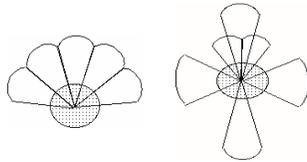


Figure 3: Homogeneous (left) and Heterogeneous(right) Sensor Suite Designs.

A homogenous sensor suite contains the sensors that have the same exact individual characteristics (max and min range, beam width); hence the only characteristics of such a sensor suite that can be varied are the number of sensors, and the placement of individual sensors. A heterogeneous sensor suite contains sensors that differ in the individual characteristics.

In this initial study, only the number of sensors and individual sensor's beam width of a homogeneous sensor suite are being evolved. The placement of the sensor is assumed to be symmetrical along the direction of flight as shown in Figure 3 (left) with the maximum sensor range of 200.0 units.

5.2 IMPLEMENTATION OF EVOLUTION

5.2.1 The Learning System

The sensor suite characteristics of the sensory payload for the MAV are evolved using GENESIS, a standard GA which maintains a "population" of candidate solutions to the objective function $f(x)$:

$$P(t) = \langle x_1(t), x_2(t), \dots, x_N(t) \rangle$$

where x_i represents a vector of parameter to the function $f(x)$ whose value is to be minimized. For each generation, the current population is evaluated using user-defined

fitness/evaluation function, and, on the basis of that evaluation, a new population of candidate solutions is formed using standard GA operations. More details about GENESIS can be found in [Grefenstette 84].

5.2.2 Representation

For this study, GENESIS' floating-point representation was used. Each chromosome described the make up of a possible sensor suite. The characteristics used to describe a sensor suite included the number of sensor in a suite (1-32), and the sensor area coverage (5-30 degrees) of the individual sensor in that suite.

5.2.3 Fitness Function

In order to fulfill the objectives of this study, each design of a sensor suite has to be evaluated based on the number of sensors in the suite and on its performance of the task. The fitness function returns a value that is to be minimized by GENESIS (see Eqn. 2).

$$\text{payoff} = \sum[(c1 * \text{num_of_sensors}) + (c2 * (1.0 / \text{MAV_performance}))] \quad (\text{Eqn. 2})$$

where $c1$ and $c2$ are constants used to weight the influence of the parameters, number of sensors in the suite and the MAV performance, on the sensor suite configuration being evolved. The MAV performance of the task is the measure of the performance of the best decision ruleset learned by SAMUEL using the sensor suite being evaluated by GENESIS. This forces SAMUEL to perform a whole learning experiment (60 generations) for every member of the population evaluated by GENESIS.

6 PRELIMINARY RESULTS

Figure 4 shows the learning curves for different designs of sensor suites where each sensor suite is defined by the number of sensor in a suite and the beam width of individual sensor. The plot shows the average performance (over 100 trials) of the best-so-far individual in the current population.

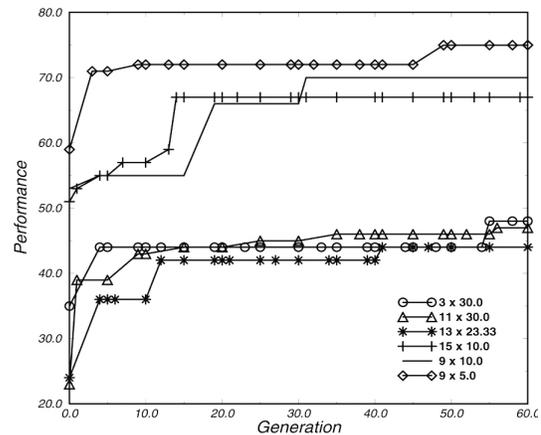


Figure 4: Learning curve for different sensor suite configurations.

There is a significant difference in performance of the task by the MAV depending on the sensor suite implemented. The sensor suites with narrower beam width of the individual sensors allow the plane to determine the position of the obstacles more precisely so the plane is able to perform its task more efficiently. Furthermore, the increase in the number of physical sensors in the suite doesn't guarantee the change in task performance. Since the sensor suites are evaluated based only on the number of the sensors in the suite and the task performance, the suites with useless sensors should be eliminated first, allowing the system to focus on determining the best individual sensor's area coverage for the task.

7 CONCLUSIONS AND FUTURE WORK

In this paper, we have presented an concept of co-evolution of form and function for the autonomous agents. The study of co-evolution of the sensor suite and the reactive control system for a micro air vehicle whose task is collision-free navigation, is currently in progress. Due to the complexity of the learning performed, only the preliminary results were available at the time of this publications, but even theses make us believe that the goal of finding the sensor suite with minimal number of sensors that guarantees an efficient performance of the task, is attainable.

Future work will include performing studies of more complex sensor suite designs including heterogeneous sensor suites in which range, beam width, and placement of the individual sensors are evolved. Such studies would require revising the evolutionary system to allow for variable length genomes on the GENESIS side.

Acknowledgments

The work reported here was supported by the Office of Naval Research.

References

[Balakrishnan 96] Balakrishnan, K. and V. Honovar. "On Sensor Evolution in Robotics." In Koza, Goldberg, Fogel, and Riolo (eds.), Proceedings of 1996 Genetic Programming Conference – GP-96; MIT Press, pp. 455-460; 1996.

[Floreano 96] Floreano, D. and F. Mondada. Evolution of Homing Navigation in a Real Mobile Robot. In IEEE Transactions on System, Man, and Cybernetics – Part B; 26(3) 396-407; 1996.

[Funes 97] Funes, P. and J. Pollack. "Computer Evolution of Buildable Objects." Husbands and Harvey (eds.), Proceedings of The Fourth European Conference on Artificial Life, MIT Press, pp.358-367; 1997.

[Grefenstette 84] Grefenstette, J. J. "The User's Guide To GENESIS." Technical Report CS-83-11, Computer Science Department, Vanderbilt University, Nashville, TN; 1984.

[Grefenstette 91] Grefenstette, J. J. "The User's Guide to SAMUEL, Version 1.3." NRL Memorandum Report 6820, Naval Research Laboratory, 1991.

[Harvey 92] Harvey, I., P. Husbands, and D. Cliff. "Issues in Evolutionary Robotics." Proceedings of the Second International Conference on Simulation of Adaptive Behaviour; MIT Press Bradford Books, 1993.

[Husbands 96] Husbands, P., G. Jermy, M. McIlhagga, and R. Ives. "Two Applications of Genetic Algorithms to Component Design." In selected papers from AISB Workshop on Evolutionary Computing; Fogarty, T (ed.), Springer-Verlog; Lecture Notes in Computer Science, 1996.

[Lee 96] Lee, Wei-Po, J. Hallam, and H. H. Lund. "A Hybrid GP/GA Approach for Co-evolving Controllers and Robot Bodies to Achieve Fitness-Specific Tasks." In Proceedings of IEEE Third International Conference on Evolutionary Computation, IEEE Press, NJ; 1996.

[Lichtensteiger 99] Lichtensteiger, L. and P. Eggenberger. "Evolving the Morphology of a Compound Eye on a Robot." Proceedings of the Third European Workshop on Advanced Mobile Robots (Eurobot '99); 1999.

[Lund 97] Lund, H. H., J. Hallam, and W-P. Lee. "Evolving Robot Morphology." Proceedings of IEEE Fourth International Conference on Evolutionary computation; IEEE Press, NJ, 1997.

[Lund 98] Lund, H. H. and O. Miglino. "Evolving and Breeding Robots." In Proceedings of First European Workshop on Evolutionary Robotics; Springer-Verlag, 1998.

[Mark 98] Mark, A., D. Polani, and Thomas Uthmann. "A Framework for Sensor Evolution in a Population of Braitenberg Vehicle-like Agents." In C. Adam, R. Belew, H. Kitno, and C. Taylor (eds.), Proceedings of Artificial Life IV; pp. 428-432; 1998.

[Nolfi 94] Nolfi, S., Floreano, D., Mighno, O., and F. Mondada. "How to Evolve Autonomous Robots: Different Approach in Evolutionary Robotics." In R. Brooks and P. Maes (eds.), Proceedings of the International Conference Artificial Live IV; MIT Press, pp.190-197; 1994.

[Sammur 92] Sammut, C., S. Hurst, D. Kedzier, and D. Michie. "Learning to Fly." In Proceedings of the Ninth International Conference on Machine Learning, Aberdeen; Morgan Kaufmann, 1992.

[Schultz 91] Schultz, A. C. "Using a Genetic Algorithm to Learn Strategies for Collision Avoidance and Local Navigation." Proceedings of the Seventh International Symposium on Unmanned Untethered Submersible Technology, University of New Hampshire Marine Systems Engineering Laboratory, pp. 213-215; 1991.

[Schultz 96] Schultz, A. C., J. J. Grefenstette, and William Adams. "RoboShepherd: Learning a complex behavior." Presented at RobotLearn96: The Robotics and Learning workshop at FLAIRS '96; 1996.

[Sims 94] Sims, K. Evolving 3D Morphology and Behavior by Competition. In R. Brooks and P. Maes (eds.), Proceedings of the International Conference Artificial Live IV; MIT Press, Cambridge, MA, pp.28-39; 1994.